DRAFT

STATEMENT OF NEED (SON)

For the

Acquisition of the NOAA Operational Central Computer System

DRAFT

Table of Contents

This Statement of Need establishes the project objectives, purpose, and the expected results for this procurement. The contract to be awarded will provide for two high performance computing systems, ancillary equipment components, system administration, support and maintenance services, including applications specialists who will assist the government in making a smooth transition to the new system.

Through this procurement, the government will acquire two high performance computer systems (Primary and Backup) to be used by the National Weather Service in performing its operational weather and climate forecasting mission for eight years. Public weather and climate guidance is issued continuously, 24 hours/day, 7 days/week, 365 days/year, on a strict, time-critical schedule, based on computer products generated on the Operational Central Computer System (OCCS) operated by the National Centers for Environmental Prediction (NCEP). One high performance computer system will be

housed in a facility provided by the contract awardee (see Appendix B). The other system will be housed in the NASA IV&V facility in Fairmont, West Virginia.
The computer systems will be used to provide primary weather and climate guidance, as well as to provide a backup capability for that mission. The current Central Computer System (CCS) is located in the IBM Gaithersburg Computer Facility.

1.0 ACQUISITION GOAL

1.1. The primary acquisition goal is to procure complete and reliable computer systems configured from commercially-available items including processors, storage, software, support and maintenance services that will allow NOAA to meet its increasing processing and strict scheduling requirements. The Primary and Backup systems will be used, on a 7x24x365 basis, to perform numerically demanding environmental simulations of the earth's atmospheric and oceanic circulation patterns, data assimilation (especially from radar satellites) and product generation in a time-critical operational environment. The system will also be used to develop advanced forecasting techniques for future operational implementation. The OCCS must perform a wide range of computational tasks including executing not only complicated prediction models and data assimilation (capability demands), but also data analysis and product generation (capacity demands) with accompanying high-volume data storage and retrieval (I/O demands). The two systems must be configured to support a seamless and rapid shift of the Production workload from one system to the other. NOAA anticipates that NCEP's capability workload, highlighted in Section 3: NOAA Workload Projection, will drive the OCCS configuration. However, efficient capacity type workload processing and an expandable storage system are also important. System configuration features, such as hardware, software or implementation, contributing towards reduced risk for NOAA operations are critically important. The current operational computer site in Gaithersburg, Maryland will not be available after the OCCS is fully operational. The vendor will be responsible for customizing both the vendor provided site and Fairmont computer sites to meet the facility requirements of the offered system and to be consistent with the requirements stated in the Statement of Need.

1.2 NOAA seeks a contract with a base period of performance of 51 months with an option for one additional 48-month period for a total period of performance of eight years and three month. Some funds should be available for program initiation in FY 2006.

1.3. NOAA expects the current CCS high performance computer, an IBM Power4 Cluster, to be saturated by the model and data assimilation enhancements scheduled for implementation in FY 2005. Therefore, the purpose for acquiring a replacement system is to increase the CCS's ability to perform these enhanced operations while meeting NOAA's time-critical needs. The guidance products created on the CCS fulfill the *NWS' Central Forecast Guidance* and *Public Warnings/Protection of Life and Property* missions. For more details on the current system configuration see *Appendix A*.

1.4. The Government is seeking a solution that considers NOAA's present computing configuration and software environment, its plans to expand computational power, and its need for operating services for a minimum fifty-one month period following contract award.

1.5. At the core of this acquisition is the operational nature of NOAA's workload, which governs the intended use of the proposed system. Because the system will support NWS' life and property protection obligations, a chief requirement is for the non-disruptive transfer of continuous daily operations and schedule-sensitive product cycles to the OCCS. NOAA requires that the OCCS meet the highest levels of reliability in order for the NWS to meet its obligations to the public. The system must also have the ability to meet NOAA's future demands through non-disruptive, periodic enhancement.

1.6. NOAA's numerical prediction models are run many times daily and are synchronized to the worldwide data collection cycle. The completed guidance products are disseminated on a fixed time schedule determined by long-standing agreements, which are inflexible. Conversely, delaying the start of model processing to collect and incorporate the most recent observational data possible enhances accuracy. Therefore, NOAA's most computationally challenging models are generally run in the smallest possible schedule "windows." Average model run start and end times (NOAA's operational "windows") are available via http://www.nco.ncep.noaa.gov/pmb/nwprod/prodstat/ . A robust capability-oriented system is required to meet this goal. NOAA's aggressive capability application growth plans require a system with substantially increasing power over the life of the contract.

1.7. Throughout the Statement of Need, "Primary System" refers to the primary computer system. The Backup System is expected to be an identical system with respect to major components such as processors, memory, interconnect fabric, disk and other elements required to achieve a balanced system suitable for meeting NOAA's computational requirements, unless otherwise stated.

2.0 OBJECTIVES

2.1. The acquisition objective is to improve the accuracy, coverage, while retaining the timeliness of NWS critical mission products by enhancing the computer processing resources used to produce them. NOAA expects an OCCS that balances its capability, capacity, I/O demands and the ability to transfer Production from the Primary to the Backup system. The contract must be flexible enough to provide NOAA with the desired period of performance, the ability to acquire and support ancillary systems (developmental platforms, servers, workstations, storage, etc.) and provide information technology services in support of NOAA's mission.

2.2. Innovation and efficient solutions are encouraged in addressing the Government's needs. The basic tenets and provisions of this Statement of Need establish the minimum acceptable performance measures based on NOAA's long-term experience in performing this mission. Newer technologies or a different approach may provide opportunities to increase performance or enhance efficiency.

2.3. The government needs:

2.3.1. **General System Attributes**

The OCCS will consist of complete and balanced high performance computing systems capable of executing NOAA's operational workload over the life of the contract. The major components are expected to consist of hardware, software, maintenance and support services. The OCCS computers shall be composed of processors, memory, interconnect fabric, disk and other elements required to achieve a balanced system suitable for meeting NOAA's computational requirements. In addition, the OCCS shall have access to sufficient disk storage and a longer term data storage system to fulfill balance requirements. All these components shall be connected via a framework supporting sufficient Input/Output rates to support NOAA's workload, including recovery from outages and supporting interfaces to NOAA's network and communication architecture. The two OCCS systems, Primary and Backup, must be supported such that the Production workload can be moved from one system to the other in a non-disruptive manner and with minimal loss of productivity. NOAA plans to operate the system housed in the vendor provided site as the primary production system but if the vendor has a more efficient plan, for example, one that includes load-balancing a portion or all the workload, the Government will entertain such a solution.

As a minimum the vendor is required to provide system that satisfies the Governments need with respect to the follow major attributes:

1. Aggregate system dependability

2. Extent and duration of any required code conversion, including Government and vendor support.
3. Numerical reproducibility
4. Integrated software engineering and development environment
5. System upgradeability and serviceability
6. Disk subsystem performance, resiliency and reliability
7. Hierarchical storage management system performance, resiliency and reliability
8. Input/Output balance between system, disk and HSM
9. Workload administration, scheduling, monitoring and execution
10. Network connectivity and performance
11. Production workload balancing with development workload
12. Support Personnel (including key personnel)
13. Training and Documentation
14. Data Migration Plan
15. Runtime variability
16. Primary/Backup system design

Technical requirements are described in the following sections. Some items specified within those sections are more critical to the overall functionality and productivity of that environment than others; therefore a three-level priority ranking has been used to describe their relative importance:

Priority Level	Description
Level 1	Items categorized as "Level 1" priority are of the most important, or critical, to satisfying the current and future needs of the CCS. Vendors are encouraged to provide all "Level 1" items.
Level 2	Items categorized as "Level 2" priority are of sufficient significance to the CCS that the effectiveness of that environment would suffer if they are not provided. Vendors are encouraged to provide as many "Level 2" items as possible.
Level 3	Items categorized as "Level 3" priority are considered to be items which enhance the utility of the system and are desirable attributes for the CCS. Vendors are encouraged to provide as many "Level 3" items as possible.

2.3.1.1. <u>System Performance</u> [Level 1]

The vendor must provide a pair of systems, Primary and Backup, that are capable of maintaining the NOAA operational product delivery schedule at NCEP, throughout the life of the contract, as the operational workload continually increases, as described in Section C3. The Government is not seeking a system that can execute an environmental application as quickly as possible. Rather the Government is seeking a system that can execute increasingly complex environmental applications within the same fixed wall-clock time window.

2.3.1.1.1 System Performance - Ancillary System [Level 3]

Vendors are encouraged to provide innovative computing solutions that may represent a significant enhancement in performance. Access to an "ancillary system" designed to provide NOAA with experience in using next generation hardware and/or software or to offer increased computational power for a particular meteorological or oceanographic application for which the primary system is not well suited, is desirable.

2.3.1.2. <u>Numerical Reproducibility/Accuracy</u> [Level 1]

The CCS must exhibit reliably consistent bit-for-bit numerical reproducibility of benchmark results. Numerical results of benchmark runs must reproduce "known truth" (NCEP control run for the same case) out to 5 decimal places or to within the prescribed precision given in the individual benchmark instructions.

2.3.1.3. <u>Run Time Variability</u> [Level 1]

The CCS must exhibit reliably consistent runtimes on successive runs of benchmark and Production codes, worst case being within 2.5% of average. The OCCS must issue automatic alerts whenever run time variability exceeds 2.5% on routine production suite jobs.

2.3.1.4. <u>System Dependability</u> [Level 1]

Both the primary and backup systems must independently satisfy the following dependability requirements, all measured monthly.

The CCS must provide sufficient available computation resources to sustain the on-time generation (within 15 minutes of target completion times) of numerical forecast products at a rate of 99.0% or better, measured on a monthly basis. The Government expects no degradation in product generation during planned, routine failover testing between the primary and backup systems. During emergency, non-scheduled failover, the Government expects no more than a 15 minute delay in the initiation of production on the backup system.

The OCCS must provide operational use time (see glossary) of at least 99%.

The OCCS as a whole must achieve a high level of Reliability, as measured by Mean-Time-Between-Failure (MTBF) of at least 240 hours (see glossary).

The OCCS shall achieve 99% system availability (see glossary).

Vendors are expected to deliver a full suite of system administration tools to achieve the required system dependability. Vendors should describe these tools and associated error detection, correction and logging attributes. Vendors should describe remote access, diagnostics and service attributes. Vendors are expected to provide key personnel to support system dependability requirements.

2.3.1.5. System Serviceability [Level 1]

The OCCS must be capable of being serviced and/or upgraded with minimal disruption or impact to NOAA Operations. This includes both system hardware and software enhancements. No scheduled downtime is provided for service. Commonly replaced components such as disks, memory models, processors, power supplies and interconnect should all be hot swappable.

2.3.1.6. Memory [Level 1]

All processors (computation, interactive, and I/O) must contain sufficient memory for their workloads. Vendors are expected to configure their systems to deliver optimal, balanced performance.

2.3.1.7 Directly Attached System Storage [Level 1]

Directly attached system storage, assumed to be disk, must support I/O transfer rates commensurate with NOAA's increasing computational demands. As of 2004, the CCS achieves read/write performance of 1.1 GB/sec. NOAA expects read/write performance to remain balanced and scale up at a rate somewhat less than the computational power rate of increase. The minimum capacity is outlined in the table below (2.3.1.8.2), and must be sufficient to support operating system requirements, handle I/O for active jobs, and have adequate space to hold staged files during job execution. File residence is typically about three days, but certain files may be kept for fifteen days.

2.3.1.8 Scalable Enterprise Storage [Level 1]

An independent, consolidated and unified storage system to provide for Archive Storage is required.

Archive Storage, supported at both the Primary and Backup facilities, holds long-term infrequently accessed data in an automated high-density storage facility such as an automated tape library. The storage hardware and management software must be completely independent from the offered high performance computer(s), but should be fully integrated with shared storage. Human interaction with removable media must be minimal, and should be limited to initial loading and the final disposal or off-site relocation. Mirrored media may be removed from the system for offline shelf storage not more often than once per month. There is no requirement for compatibility with NOAA's present archive technology.

2.3.1.8.1 Hierarchical Storage Management (HSM) [Level 1]

Hierarchical migration of data across the all categories of storage must be fully automated and require minimal or no special actions by the computer user. However, the system must provide a means for the user to influence or control the process at the discretion of NOAA management.

2.3.1.8.2 Required Features [Level 1]

Scalable storage architecture having the following features:

<u>Connectivity:</u> Attachment through high speed networks to a heterogeneous collection of computers and operating systems
<u>Capacity:</u> Capacity sharing and non-disruptive expansion
<u>Performance:</u> Retrieval, staging, I/O speed, and migration across storage systems commensurate with NOAA's growth of computational power and NOAA's goal of a balanced multi-computer central computing facility
<u>Reliability:</u> Zero loss of data
<u>Availability:</u> Uptime of 99 percent
<u>Resilience:</u> Self-healing RAID technology and hot-swap maintenance
<u>Recoverability:</u> Rapid restoration of corrupted metadata independently of data

In addition, the storage system should have the following attributes:

i. Centralized management and control of all storage
ii. System-wide data security using adjustable user and group permissions or better
iii. Non-disruptive maintenance, expansion, technical refresh, and migration of data

iv. Non-disruptive background storage management functions
v. Real-time data duplication or creation of backups for important production data
vi. Time and event-triggered (point-in-time) control of the following storage system tasks:
vii. Dynamic priority shifting to adjust performance during operational catch-up
viii. Threshold activated archiving
ix. Metadata capture and backup or journaling
x. Remote mirroring to a distant processing site
xi. Proven supplier support for enterprise applications in a heterogeneous environment
xii. See table below to determine trend of minimum usable (formatted) storage capacities in line with NOAA's estimate of future requirements. NOAA expects disk and tape upgrades to coincide with computational power upgrades.

Fiscal Year	System Storage (TB)	Archive Storage (TB)
2004	49	2500
2005	100	3000
2006	210	10000
2007	315	15000
2008	420	20000
2009	630	30000
2010	840	40000
2011	1260	60000

The 2004 estimates reflect available disk and tape on the Primary and Backup systems

2.3.1.8.3 Desirable Features [Level 2]

Transparent transition to new technologies by legacy applications
Controller to device ratios in line with optimal performance goals
Capacity sharing by reassignment of allocation units
Data sharing, integrity interlocking, and synchronization
Dynamic multi-system allocation of serial access, removable media devices
Robust set of storage management and control tools
Workload analyzer for performance guidance (including summaries by application)
Ability to rebalance contention by re-mapping logical to physical units
Division of storage pools by group and project
Extensive reporting and trending capabilities

Usage tracking by user, account number, project, division, or branch
Support for service-level agreements and priority of service commitments
Support for integrated storage clustering

2.3.1.8.4 Option for re-purposing GFE after the end of the current contract

At the end of the current period of performance, September 30, 2006,NCEP-owned peripheral equipment at both Gaithersburg and Fairmont will become available for re-utilization or trade-in towards meeting the storage requirements of this contract Equipment which will become available at that time will consist of:

STK Powderhorn silos (2)/9940B tape drives (16) Gaithersburg
STK Powderhorn silo/9940B tape drives (4) Fairmont

Offerors who plan to use this equipment or trade for value must indicate so in their response to the RFP. An implementation plan of a few paragraphs should be included.

2.3.1.8.5 Data transfer to the Scalable Enterprise Storage subsystem

Offerors shall provide a plan describing how data will be migrated from the current computer systems in Gaithersburg, Maryland and Fairmont, WV to the OCCS Scalable Enterprise Storage subsystem archives. This will be a vendor responsibility with Government oversight. Most data is binary numerical fields, with some source code files. There are no commercial databases to be converted. The task will consist primarily of organizing and cataloging data for future retrieval from a diversified collection of computers, including the OCCS.

2.3.1.9. Operating System [Level 1]

The OCCS must be a UNIX or Linux based system which supports both the C shell (csh) and Korn shell (ksh). Vendors should report conforming standards.

2.3.1.9.1 Operating System [Level 3]

Vendors should provide an opportunity for the Government to pre-test new releases on an OCCS test system.

2.3.1.9.2 Accounting [Level 1]
The system must be delivered with an accounting package capable at a minimum of tracking all aspects of system utilization by individual account.

2.3.1.9.3 Security [Level 1]

The Vendor must provide an operating system which will provide adequate security features to ensure integrity of NOAA's critical mission.

2.3.1.10. Processors [Level 1]

The OCCS must be provided with processors that support both 32-bit and 64-bit fixed and floating-point numerical representations. Vendors should report conforming standards.

2.3.1.11. Single Point System Administration [Level 1]

The OCCS must have the capability of being administered via a single stand-alone workstation, on-site and/or remotely, capable of administering either the entire system simultaneously or on a node-by-node-basis.

2.3.1.12. Network Connectivity [Level 1]

The OCCS shall be capable of supporting connectivity to external networks using either an ATM OC-12 (622 Mbps), Gigabit Ethernet (1000 Mbps), or other standard interfaces having the equivalent bandwidth. The need for interface bandwidth is expected to scale logarithmically to an OC-48 level over the term of the contract. The OCCS network interfaces shall be fully redundant to assure high availability.

2.3.2. **Software**

2.3.2.1. Compilers [Level 1]

As a minimum, the OCCS must be provided with compilers that support programs written in FORTRAN, C, and C++. Vendors must report conforming standards.

2.3.2.2 Parallel Programming and Debugger [Level 1]

The OCCS must support MPI, MPI-2 I/O, OpenMP and a hybrid programming model. The OCCS should be provided with TotalView™ or a functionally equivalent OpenMP and MPI parallel application debugger. Vendors must report conforming standards.

2.3.2.3. Parallel Programming and Debugger [Level 1]

Full MPI-2 support is desirable.

2.3.2.4. Serial Debugger [Level 1]

The OCCS must be provided with the standard UNIX **dbx** debugger and/or a functionally equivalent product for debugging serial codes.

2.3.2.5. Libraries [Level 1]

The OCCS must be provided with X11R6, Optimized Scientific Subroutine Library (OSSL), Fast Fourier Transforms (FFT).

Vendors must provide subroutine libraries, especially mathematical & statistical.

2.3.2.6 File system [Level 1]

The OCCS must have a high performance, resilient common file system that provides isometric performance and accessibility from all nodes of the system. Vendors should fully describe the file system(s) available including features such as aggregate and single-file I/O performance, performance tuning and administration capabilities. Capabilities for problem solving should be detailed.

The size of the file system and the maximum number of files and maximum file size should be reported.

2.3.3. **Support**

2.3.3.1. System Administration [Level 1]

The OCCS must provide sufficient support staff to administer the computation system and all peripheral devices and to work with NOAA staff members who help support operations and R&D work. NOAA anticipates that the successful offeror will provide a system administration staff of at least six persons to administer the two systems, based on past experience.

System administration support is required 7x24x365. Typically, on-site staff work normal business hours. On-call and emergency staffing arrangements must be sufficient to meet NOAA's production requirements.

2.3.3.2. Application Support [Level 1]

The successful offeror must provide sufficient support staff to assist the Government staff with code optimization, data migration, training, and code conversion. The Government anticipates the successful offeror will provide an on-site application support staff of at least two persons, based on past experience. The Government expects to have full-time access to key vendor personnel involved in benchmark run preparation for a period of at least 90 days following the acceptance of the primary OCCS

2.3.3.2.1 Application Support [Level 3]

Vendors should provide opportunities for collaborative work to develop innovative support functions.

2.3.4. **Job Scheduling**

2.3.4.1. Checkpoint Restart [Level 2]

The OCCS must include the capability to checkpoint a running process, via administrator or software action, then restart its execution from the checkpoint start. This capability should be provided for serial and MPI and OpenMP applications.

2.3.4.1.1 System Initiated Checkpoint Restart [Level 2]

System initiated checkpoint restart is a desirable attribute in special circumstances. Vendors should use the benchmark codes to determine the system resources required to support system initiated checkpoint restart and estimate the time required to checkpoint the benchmark codes on a full system.

2.3.4.2. Workload Management Software [Level 1]

The OCCS must include software capable of efficient pre-emptive priority and Gang scheduling. Workload management software must function dynamically (non-disruptive) and be flexible enough to meet NOAA's operational requirements including numerous scheduling changes to accommodate severe weather events.

2.3.5. **Periodic Performance Improvements** [Level 1]

The OCCS must include provisions for periodic performance improvements after the initial system delivery, through technology refreshment reengineering offerings. These periodic enhancements should occur approximately every 24 months, through both the base period and option period at no increase in cost. Enhancements should be timed to best meet the Government's needs for increased performance and reliability. The need for a balanced system (processing, I/O, storage, bandwidth) is inherent in these requirements. The Government may opt to accept, defer, or reject these offerings depending on a cost-benefit analysis performed for each proposed upgrade. The government also reserves the right to improve performance through sources outside this acquisition.

2.3.6 **Workload Migration** [Level 1]

Even though the benchmark codes represent the most challenging codes with respect to computational resource requirements and code complexity, they represent only a small fraction of the total number of codes running in the NCEP Production Suite. The Production Suite is currently comprised of 800 codes, 20 of which utilize MPI. The Government will be responsible for the migration of the 800 applications with assistance from the vendor. Depending on the architecture offered and the complexity of the migration, the vendor will be expected to provide the appropriate level of support and training during the migration. The Government expects to have full-time access to key vendor personnel involved in benchmark run preparation for a period of at least 90 days following the acceptance of the primary OCCS.

2.3.7 **Pre-delivery System** [Level 1]

The government needs pre-delivery access to a system with characteristics similar to the proposed system within one month after contract award. This access will be used to develop codes and test converted work for the new system during normal business hours. This access must continue until the acceptance of the Primary OCCS.

2.3.8 **System Installation** [Level 1]

System installation and/or site preparation work is presently scheduled to begin not later than October 1, 2005. System Acceptance (Section E.2.2) is required not later than July 1, 2006.

2.3.9 **Reserved**

2.3.10 **Training** [Level 1]

The successful offeror must deliver a comprehensive training program for OCCS administrators and end users. The program must include periodic refresher training for new releases. Classes for end users should be on-site, custom designed for NOAA's scientific staff and should include classes on system basics, compiler features, parallelization techniques, performance tuning and HSM features. Meteorological and Oceanographic applications and associated libraries should be highlighted, if applicable. Training for system administrators should be a combination of standard classes taught at the vendor's facilities and on-site custom classes. Pertinent topics including, but not limited to, advanced system administration, file systems, resource management, performance tuning, system troubleshooting and failure analysis should be offered.

2.3.11 **Documentation** [Level 1]

Vendors are expected to deliver extensive, on-line system documentation. Vendors must specify what printed manuals are available and which and in what quantity they propose to deliver.
Documentation should include but should not be limited to;

Site preparation and installation guide,
System hardware manual(s)
System operations manual(s)
User's manuals (system, commands, compilers, assembler, debugger(s), libraries, performance utilities, and others
System programmer's manuals
System administrator's manuals

Standard UNIX "man pages," a low level user's guide and a problem reporting procedure are required. Advanced level documentation, in support of training activities (2.3.10) are required.

2.3.12 **Earth System Modeling Framework (ESMF)** [Level 1]

Prior to the end of the option period of the OCCS contract the Government expects model system software on the OCCS will become compliant with the Earth System Modeling Framework. More information on ESMF can be found at http://www.esmf.ucar.edu/.

2.3.13 **Computer Facility** [Level 2]

The Government requires the contractor to include as a Government-exercisable option, a facility to accommodate the needs and specifications of the offered equipment (*See Appendix B*). If the proposed site is outside of the immediate Washington, D.C. metropolitan area, the contractor shall provide telecommunications between the National Weather Service Telecommunications Gateway, NCEP's offices (currently at the World Weather Building in Camp Springs, but expected to move to College Park in 2008), and the contractor's facility.

3.0 NOAA WORKLOAD PROJECTION

3.1 Computational workload projections are derived from the current NCEP model suite as it is implemented on the IBM CCS system. Table 1, below, lists the major components of the NCEP model suite along with their execution windows and computational power requirements, which are expressed as a fraction of the IBM CCS resources. Each model is extrapolated out through the 2010 time frame reflecting the increased load due to higher resolution, more sophisticated physics or additional ensemble members. Detailed model plans beyond 2010 are not presently available, but it is the government's expectation that a Moore's Law extrapolation would apply beyond that point. The projected computational requirements are expressed as multiples of the baseline IBM CCS (phase 2 upgrade) power.

3.1.1 Each model in the production suite runs in a fixed time window. In general, models may not start early (input data are not available) nor may they finish late due to strict product delivery commitments and subsequent run dependencies. These circumstances produce very high peak computational loads and drive the platform requirements accordingly. Figure 1, below—the 'jigsaw puzzle'—illustrates this complex mesh of individual model execution times in a production cycle which is repeated 4 times per day.

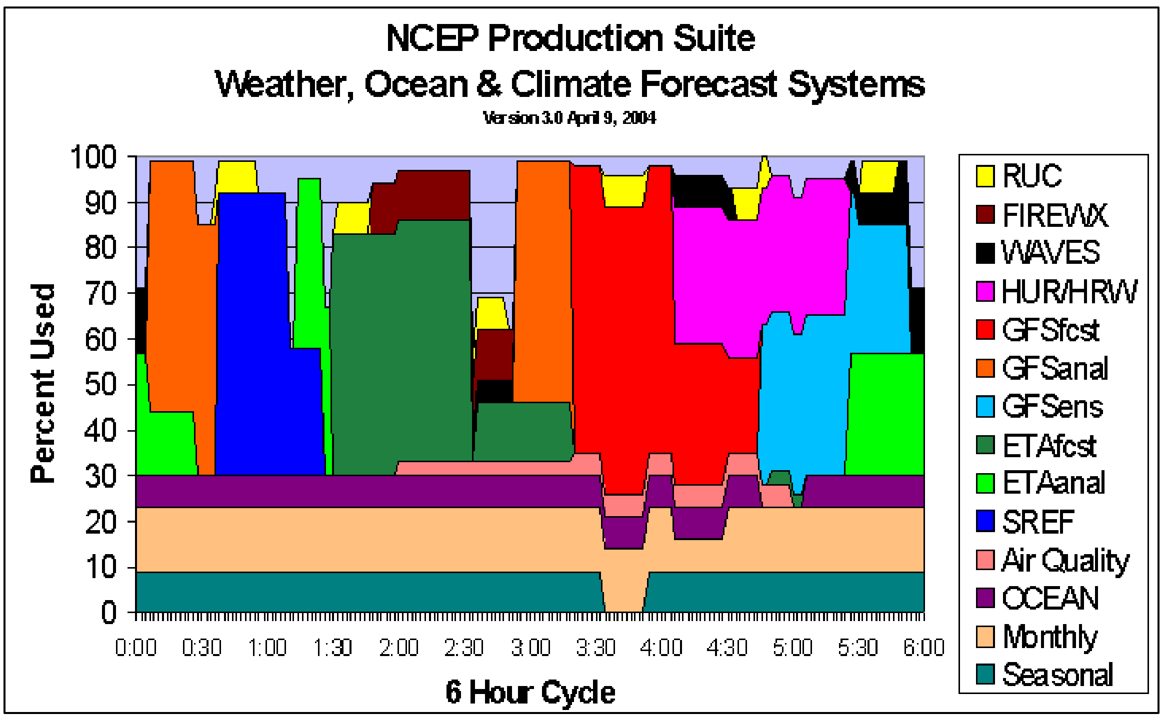

Figure 1. NCEP ¼ - day Model Cycle

Model Implementation Year and Computational Requirement					
Operational Product Model Type	6 Hr Window	Oct-04 Baseline Resolution	Oct-06 Projection - 4 X Resolution	Apr-09 Projection - 12 X Resolution	Oct-11 Projection Resolution
GFS - Global Forecast System Prediction Model	63% - 50 min 31% - 24 min 21% - 16 min	T-254 / L64	45 km / L64	40 km / L80	35 km / L100
Analysis	53% - 26 min	55 km GSI	45 km 3D Background Error Covariance	40 km	35 km 4DDA
Data Assimilation	55% - 32 min				
Ensemble Forecasts GFS - Ensemble	35% - 45 min 28% - 25 min	T126 / L42 - 15 M	90 km / L42 - 50 M	80 km / L64 - 50 M	70 km / L64 - 50 M
Short Range Ensemble	62% - 32 min 28% - 16 min	32 km - 15 M	22 km - 20 WRF M	15 km - 20 WRF M	12 km - 25 WRF M
North American Meso Guidance System Prediction Model	53% - 65 min 13% - 44 min	12 km 60 Levels	10 km 60 Level WRF	8 km 70 Level WRF	6.5 km 85 Level WRF
Analysis	37% - 16 min	12 km 3DVAR	10 km GSI	8 km	6.5 km 4DDA
Data Assimilation	27% - 40 min 14% - 20 min				
High Resolution Window Fire Wx & OCER	28% - 80 min	8 km WRF - 6 M	7 km WRF 8 M	6 km WRF 10 M	5 km WRF 10 M
Fire Weather IMET Support	11% - 66 min	8 km WRF-NMM	6.5 km nested WRF	5.5 km NAM-WRF	4.5 km NAM-WRF
Hurricane Model	30% - 80 min	18 km + 55 km L42	12 km + 40 km nests L64	8 km + 30 km nests L64	5 km + 20 km L100
Rapid Refresh	7% - 60 min 7% - 48 min	20 km L50 RUC	13 km L60 RUC	10 km L60 RR WRF	8 km L70 RR WRF
Air Quality	3% - 80 min 5% - 86 min	12 km	10 km	8 km	6.5 km
Windwaves Oceans & Coastal Zones	14% - 20 min 7% - 22 min	Global 1°, Reg 1/4°	Multi-scale Global 1/2° Global ensemble 1° 15 M	Multi-scale Global w/costal focus Global ensemble 1° 50 M	Multi-scale Global 1/3° Global ensemble 3/4° 50 M
Great Lakes	5% - 20 min		4 km	Multi-scale model	Coupled modeling (Wavewatch - WRF)
Ocean Model (HYCOM) Oceans & Shelf	7% - 230 min		Global 1/3° North Atlantic 1/12° equivalent	Global 1/6° +North Pacific 1/12° equivalent	Coupled Ocean-Atmosphere small ensemble
Coastal & Surge	7% - 110 min		East Coast High Res. Surge	Gulf of Mexico, Small ensemble	Extend to large ensemble
Climate / Coupled Model Seasonal Forecast (monthly)	9% - 340 min	T62 L64 atmosphere 100 km L40 ocean	T62 L64 atmosphere 100 km L50 ocean 2 M / day	T126 L64 atmosphere 40 km L50 (new) ocean 2 M / day	T126 L64 atmosphere 40 km L50 (imp) ocean 2 M
Monthly Forecast (weekly)	14% - 336 min 7% - 24 min		T126 L64 (improved physics) 40 km L50 2-4 M / day	T126 L64 (atmosphere) 40 km L50 4 M / day	T170 L64 (imp physics) 35 km L60 4 M / day
Ancillary Processing Load	5% - Continuous				

Table 1. Computational Requirements Projection as Multiples of CCS Phase 2 power.

3.1.2 The government expects that commercially available computational power will follow an exponentially decreasing cost curve (as dictated by Moore's Law). Figure 2, below, is a general projection of computational power that should be available at the fixed OCCS budget level. This projection is anchored at the point on the exponential growth curve corresponding to the Phase 2 CCS machine's power and cost.

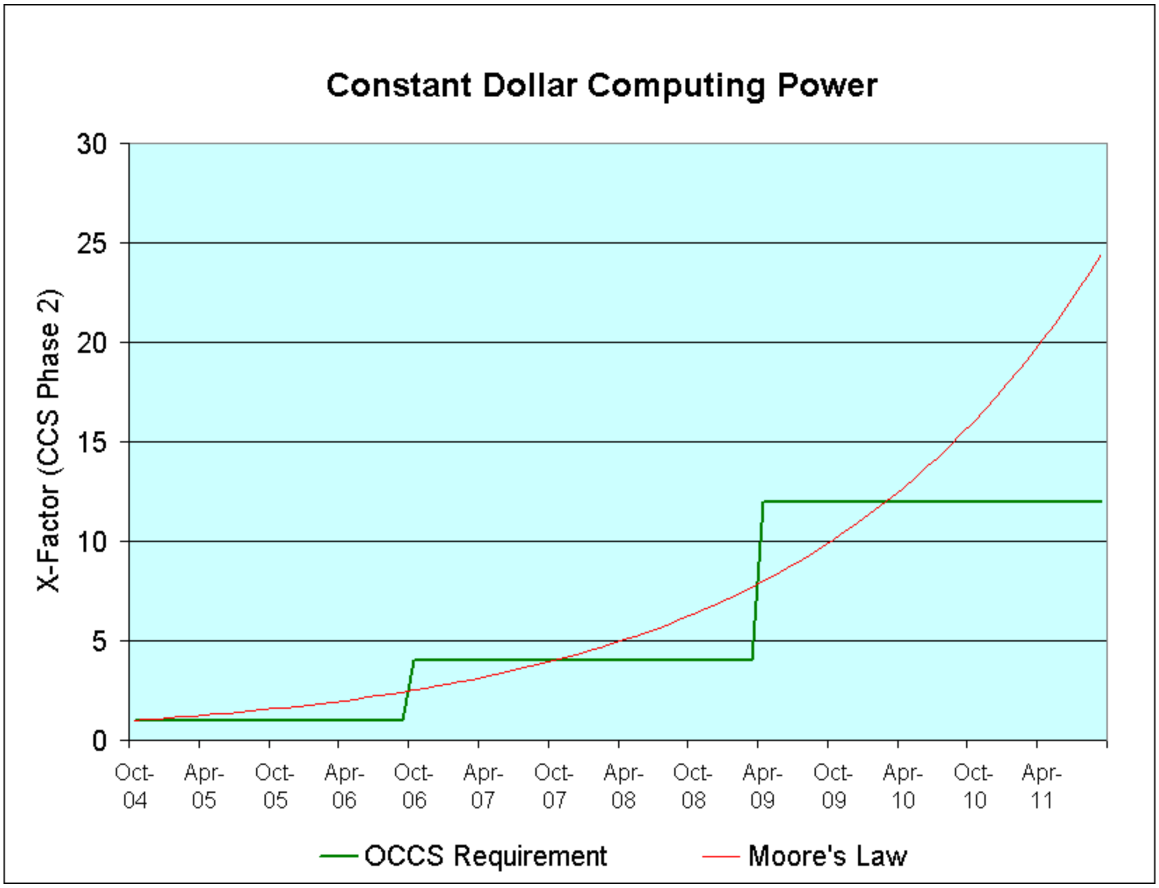

Figure 2. NCEP HPC Projection

3.2 The computer system will be part of an evolving heterogeneous computing environment spanning multiple network-attached systems. There are geographically distributed interests who will access the computer system to obtain data, use interactive display software and databases, and, in turn, supply value-added products for use by other center clients. Although evolution and design of this broader computing environment is to be guided by NOAA's own network planners, the Government may need additional integration expertise and services relative to attaching the new computer system to the existing infrastructure.

Glossary

Operational Use Time: Operational use time shall be defined as that period of time in which the batch computation and spare processors, plus all necessary I/O, server, networking and storage subsystems, are available for execution of the production suite.

Reliability: Reliability will be expressed in terms of Mean-Time-Between-Failure (MTBF). A failure is defined as any scheduled or unscheduled event that triggers down time. Individual components of the system can fail and not be counted as a system failure if the system has sufficient redundancy, error correction or resiliency to allow the failure to not affect individual jobs running in production or development.

Availability: Availability shall be determined by computing the ratio of total batch computation processor hours available for execution of the operational and developmental suites to the total batch computation processor hours each month. Time a batch computation processor is available for execution is determined by subtracting processor down time from wall-clock time. Spare nodes or processors can be configured into the batch computation processor pool upon failure of a node and/or processor in the batch computation pool to reduce down time, but down time accumulates until the spare is made available in the system for execution. Failure(s) of other components of the OCCS that result in batch computation processor(s) being unable to run jobs is equivalent to failure of those batch computation processors themselves.

Down Time: Down time is that period of time when the OCCS or a component thereof is inoperative. Down time shall commence at the time when NOAA contacts the Vendor's maintenance representative at the designated point of contact or the Vendor's answering service or other continuous telephone coverage provided to permit NOAA to make such contact to report a failure. Down time shall end when the OCCS is returned to NOAA in operable condition. The OCCS, or individual component thereof, may be declared inoperative while problem diagnosis takes place. Down time includes any time required for operating software regeneration, reinstallation, or reconfiguration. During a period of down time, NOAA may continue to use operable components of the OCCS when such action does not interfere with maintenance of the inoperable components of the OCCS

Null Time: NOAA acknowledges that, during the acceptance test period, unanticipated events may occur that are beyond the control of either the Vendor or NOAA. In such events, NOAA’s COTR and the Vendor may agree that one or more periods of time be counted as "null time". Null time does not affect utilization, availability, operational use or reliability metrics and simply extends the duration of the acceptance test and evaluation periods by the equivalent wall-clock amount (i.e. null time is as if the wall clock stopped for that period of time).

Appendix A

Configuration of NCEP's present computer system as of October 1, 2004

This information is provided solely as a reference for orientation by potential offerors. It does not reflect any requirements of the OCCS RFP, nor does it represent equipment or software preferences, functional capabilities, compatibilities, capacities, etc.

NCEP currently runs two IBM Power4 Cluster 1600 systems. One is located at IBM's Gaithersburg, MD facility. The other system, a functional clone, is located at NASA's IVV facility in Fairmont, WV. The Backup typically handles a research and development workload but is designed to handle Primary NWS operations under a continuity-of-operations concept.

Under NCEP's CCS contract, the final upgrade during the base period of performance, termed "Phase-II," was installed during July 2004 (at Gaithersburg) and August 2004 (at Fairmont). The Primary and Backup systems are expected to become operational on January 18, 2005

The two discrete computers are managed independently such that the Primary runs operational (time-critical) work and the Backup runs primarily development work but is immediately available for business continuity backup functions. They are functionally symmetrical; the primary difference being the type of cabling used for the High Performance Switch fabric. The configuration as presented below describes the combined two-computer system.

A general estimate of theoretical peak performance of NCEP's IBM Power4 system is about 8.7 teraflops.

The compute node technology throughout is IBM's Power-4Regatta nodes running at 1.7 GHz and configured with eight CPUs per node, each node having 16 MB of memory. Other highlights include:

1. 176 nodes total (160 compute nodes)
2. 1,408 total processors (1,280 compute processors)
3. 16 GB memory per node (2 GB per processor)
4. two 34.6 GB internal disk drives in each node for operating system use (and hot spare)
5. IBM High Performance Switch fabric (Federation)
6. 26.5 TB disk subsystem in Gaithersburg; 22.5 TB disk subsystem in Fairmont; all fiber channel disks in RAID 5 configuration
7. 2 PB robotic tape system in Gaithersburg (two StorageTek Powderhorn silos)

8 1 PB robotic tape system in Fairmont (one StorageTek Powderhorn silo)
9 Software levels:
    1 AIX 5.2 operating system
    2 LoadLeveler (job scheduler) Release 3.2
    3 POE (Parallel Operating Environment) Version 4.1
    4 Cluster Systems Management (CSM) Version 1.3
    5 XL Fortran version 8.11
    6 Third party commercial software
        TotalView (Parallel Debugger)
        Vampir (Parallel Program MPI Optimizer)
10 System nodes (per each computer)
    1 8 GPFS (General Parallel File System)
    2 4 Interactive nodes
11 Hierarchical Storage Management system using HPSS Version 4.5 for direct access to tape
12 20 StorageTek 9940B high capacity drives

In NCEP's configuration, only 8 nodes on each computer perform all disk I/O (excluding the OS native disk in each node). Processing nodes transmit and receive data from servers over the switch fabric

Appendix B

Option to supply an operating site for the OCCS computer

Offerors are required to include a separately priced option to provide an installation site for one of the offered systems. Funding for the new facility will be taken from the present site operating budget and other sources.

This appendix summarizes some of the characteristics required and desired by the government. Representative criteria for assessing a high quality, zero-outage facility are presented; however, the government makes no commitment to use these or any other criteria as evaluation factors. They are provided solely as an indicator of how the government may determine the quality and reliability of service provided by a facility offered under this option. Also, the government will consider feedback from prospective offerors as to alternative characteristics or assessment methods up until the draft RFP is finalized and formally issued.

The government is keenly aware of quality differences in computer space offerings and in facility support equipment. The government also recognizes that poor management of a given facility can easily offset the advantages of the finest high quality equipment and support infrastructure. The government will consider factors that substantiate the offeror's guaranteed availability, and to distinguish quality differences.

BACKGROUND:
NCEP houses its current MPP operational computer in space leased from IBM in Gaithersburg, Maryland. The computer site provides no space for NCEP staff and consequently runs as a "dark-room" operation remotely controlled by NCEP operations staff located within the NOAA Science Center (a.k.a. World Weather Building).

All building support infrastructure including physical plant operation, redundant power and cooling testing, fire protection systems, security services, recurring and emergency maintenance, capacity management, alterations, design, and inspections, etc. are provided by IBM.

NCEP's experience in a "turn-key" service-based high-availability computer facility of this caliber has been very good. NCEP's core business model does not include computer center management and facility operations. NCEP would like to continue locate it's high performance operational computer system in a purpose-built, modern, reliable computer facility.

Computer Facility Assessment Guide for a vendor supplied site:

Downtime:

Computer downtime resulting from any sort of facility-induced interruption, disturbance or outage will also be incorporated into the overall computer system downtime assessments. As a constituent part of the computer's guaranteed availability, facility downtime would be added to downtime resulting from other causes (e.g., computer hardware and software). Facility downtime begins with the first interruption to system operation that reduces government use of the system until full system functionality and usability are restored. Additional interruptions within twenty-four hours of the contractor's last notification to the government that the system has returned to service will be deemed a single extended outage beginning with the time of first interruption. Downtime reporting will be rounded down to the nearest minute.

Because high-availability is an integral component of facility design and operation, activities that negatively impact on the government's use or access to portions or all of the system include emergency maintenance, planned maintenance, deferred remedial work, pre-scheduled testing, alterations, and similar work. Consequently, planned activities will be counted as downtime unless specifically excluded by the government in advance of the work.

Risk:

The government expects zero facility downtime. In NCEP's near real-time operational processing, every interruption results in extended recovery time for production work–resulting in late products–that worsens with the outage duration by increasing the catch-up time. Offerors must provide a guarantee of sustained reliability they believe they will achieve, and a guaranteed maximum number of computer interruptions annually. The government expects not more than three interruptions annually and reliability due to facility causes not less than 0.99977 (2 hours of downtime annually).

Design resiliency:

Offerors must provide documentation sufficient for the government to determine that no single point of failure exists in any of the physical plant and facility support equipment. Offerors must identify the power conditioning, backup diesel generators, and other redundant backup systems such as multiple sources for power and cooling. Descriptions of the automated control systems, rapid switchover logic, monitoring sensors, etc. are also required.

Design compliance:

A short notation as to what standards were used by the A&E in the design of the computer facility and major support components (e.g., design of fail-safe power and cooling distribution systems). For example, list which professional organization's data center design principles, practices, and conventions were endorsed and used.

Security:

Offerors must describe what controls are in place for access security and protection from extremist groups. Written security and policing procedures should be enumerated and synopsized, along with practice drill frequency and post-drill assessment methods.

Fire Protection:

Briefly describe how the fire detection and suppression systems meet modern design principles for high-availability computer centers. Also briefly describe the disaster response and recovery plans.

Cost:

Cost data shall be presented in spreadsheet form, monthly and by incremental enhancement. If the offer consists of an upgrade that consumes additional space, then the spreadsheet should reflect the month when that space will be needed, the total number of square feet including the addition, and the square-foot rate.

All primary cost elements should be listed independently, and then summarized down to a single square-foot cost. Variables such as power and cooling requirements should be included in this square-foot rate. Supplemental costs that may vary over time must be identified and estimated. All reimbursable costs and liabilities must also be presented with the cost data.

Assessment:

The following documents shall be provided to the government. This material will be considered prior to a site visit, when the government will inspect additional documentation. For new facilities where existing documentation has not been established, an example of the offeror's planned documentation shall be provided, and clearly marked as "Planning Example".

.

REQUIRED DOCUMENTATION:

1. A guarantee of the maximum number of facility incidents and annual outage time
2. A statement documenting any single points of failure in computer conditioning
3. A description of how NCEP remote computer operators will be informed of deteriorating facility conditions such as rising room temperatures or an air handler failure
4. A copy of the site operating plan
5. A copy of the security procedures
6. A copy of the disaster recovery plan
7. One-line (logic) diagrams of the electrical service and cooling service
8. An energy density (watts per square foot) projection plotted over the contract life
9. A spreadsheet listing the type and age of facility equipment to be used. Examples are: UPS systems and power conditioners, chillers, heat exchangers, air handlers.
10. A copy of the contract statement of work for any commercial facility management company used, or the equivalent if preformed in-house. PM schedules, proactive inspections, and quality assurance methods are examples.

11 A brief (2-3 paragraphs) description of the procedures used to acquire off-site emergency service including minimum response times and escalation procedures
12 A statement (one paragraph) as to how coverage and services are made available after-hours and on holidays
13 A statement (paragraph) projecting the minimum UPS power protection period (survival time) when utility power fails. The worst-case instance covers full functionality of the system(s), and over the entire period of performance (i.e., should load vary during the contract, or battery performance decline due to age).
14 A brief description of the fire protection systems and certification standards
15 A list of design standards complied with by the facility architect (see above)
16 A brief description of automated facility controls such as computer-managed fail-over systems
17 A brief description of quality assurance control processes for defect corrections
18 A brief description of the tools used by management to track performance of service level commitments and to manage the facility (Examples: CAD programs, equipment inventory management software, capacity planning and tracking, etc.)
19 A bio (curriculum vitae) of the facility manager's experience and training
20 A description of facility alterations and changes to be made to the offered space if the offeror is successful
21 Cost data as described above

DURING THE SITE VISIT, THE GOVERNMENT WILL INSPECT:

I. Mechanical rooms
II. Raised floor plenums; drainage; cable tracks, labeling, and management
III. Network cabling protection and redundancy provisions
IV. Power and cooling distribution and control systems
V. Logs of equipment failures, corrective actions taken and maintenance results
VI. Preventive maintenance schedules
VII. Emergency plans, and results of drills
VIII. Testing procedures and schedules
IX. Training provided to facility managers, maintenance and security personnel
X. Safety and fire protection equipment and operation
XI. Any other relevant materials that will enable the government to assess the reliability and safety of the facility